Subscribe Login Home Archive Contribute Careers About

Blood flow in the brain, multi-scale modeling, and more — life as an early-career HPC researcher

iSGTW speaks to Derek Groen, a post-doctoral researcher from the Centre for Computational Science at University College London (UCL), UK. He'll be presenting his work into the optimization of hemodynamics simulation code at ISC'14, and he tells iSGTW why the event is not to be missed by early-career researchers.



Posted on APR 16 2014 7:11AM



"I'm perhaps a bit unusual in having worked with astrophysics, blood flow, and materials applications, but the computing ties everything together." Image courtesy Derek Groen.

iSGTW speaks to Derek Groen, a post-doctoral researcher from the Centre for Computational Scienceat University College London (UCL), UK. He'll be presenting his work into the optimization of hemodynamics simulation code at ISC'14,and he tells iSGTW why the event is not to be missed by early-career researchers.

At ISC'14, you're going to be giving a presentation in the session on advanced re-engineering of high-performance computing (HPC) applications. Could you tell the *iSGTW* readers a little about this?

I'm going to present optimization work that we've done for a high-performance hemodynamics code named HemeLB. It's used to model blood flow in blood vessels in the brain and I'll be talking about work we've done to improve this code.



Andrew Purcell European editor

Share this story

f

8⁺ in

t과 Republish

Tags

Why is hemodynamics simulation important?

We want to understand blood flow in the brain. It can be very difficult for clinicians to measure things like flow velocity, especially under conditions which may be hard to replicate in a clinical setting, such as when the subject undergoes a change in heart rate or blood pressure. Using the HemeLB code, we're able to make predictions for flow conditions that are different to those easily observed by clinicians in the hospital. This is where the supercomputers come in: it is not uncommon for us to require at least a few hundred, or even tens of thousands, of cores.

I've read that you're going to be involved in the CER2EBRAL project. What exactly is this?

The CER2EBRAL project, which is funded by the Qatar National Research Fund for three years, will be launching just a few months from now. Nevertheless, we've already done some preliminary work and have prepared a conference paper. The idea with the CER2EBRAL project is to create an automatic or at least semi-automatic - clinical workflow based on the HemeLB code, which aims to help clinicians to predict the outcomes of vascular surgery in the brain. At the moment, we run HemeLB on large supercomputers, such as ARCHER, and resources from the European Grid Infrastructure (EGI). So, we need to make HemeLB more suitable for execution within a clinical environment. We intend to optimize the code for running on many-core architectures and then try it out within a clinical environment on a GPU-enabled machine. We'll also be working on improving the efficiency and accuracy of how 3D models are produced from clinical angiography scans.

I believe you're also an associate fellow of 2020 Science. What are the aims of this project?

2020 Science is a large British program, which funds and supports a large number of fellows within life sciences research. These fellows and associate fellows such as myself - all coordinate and interact with one another to develop common visions for the future of computational research. A good example of this is a paper we recently published called 'Ten Simple Rules for Effective Computational Research'.

2020 Science also aims to establish strong interdisciplinary collaborations. Computational skills are often relatively isolated within specific scientific disciplines, so by interconnecting these disciplines and creating collaborative research projects, we can help prevent researchers from reinventing the wheel.

On the subject of interdisciplinarity, I can see from your publication history that you've been involved in work with a number of scientific fields. Is this a common situation for researchers with expertise in HPC-related techniques?

I'm perhaps a bit unusual in having worked with astrophysics, blood flow, and materials applications, but the computing ties everything together. My expertise in distributed computing and parallelization has allowed me to work in different application fields. Although the scientific problems can be very different, the underlying methodologies used are often very similar. For example, with both cosmological simulations and molecular dynamics simulations, one is working with potentials and interacting particles. There are lots of things that one can port from one discipline to another.

Of course, it's important not to overplay the similarities. When I was working on the MAPPER project, I did a survey of different communities working on multi-scale modelling and simulation and I saw quite a lot of differences in terms of the approaches adopted within different scientific disciplines.

MAPPER stands for 'Multiscale Applications on European e-Infrastructures', but what exactly does the project entail?

MAPPER has been actively involved with both EGI and the Partnership for Advanced Computing in Europe (PRACE), because one of the main goals of the project was to bring distributed multi-scale simulations to production e-infrastructures like these. We also created a number of multi-scale applications using existing codes and existing tools. These were then demonstrated and used on EGI and PRACE resources.

In addition, we've done research to ascertain the extent to which the European e-infrastructure research landscape is already suited to the needs of multi-scale modelling, both in terms of technology and policy. One outcome of this was that we set up facilities on several EGI resources for advanced reservation, for example in PL-Grid, which is crucial if you have individual codes running on different machines at the same time. There are now also PRACE resources that offer similar advanced reservation, too.

In addition, I believe you're involved with a skills-sharing initiative known as

'software carpentry' through the Software Sustainability Institute. Could you perhaps tell the *iSGTW* readers more about this?

I recently received a fellowship grant from the Software Sustainability Institute, through which I'm able to set up and attend a number of workshops. In the past, I've given a few talks about coupling and parallelizing codes and I've recently volunteered to do some teaching on topics related to software carpentry, particularly in distributed and parallel computing. If people want to learn more about how to parallelize their code, how to run their code in a distributed environment, or something similar, I am able to visit them for a day or two and teach them about these things. Software carpentry is important because although quite a lot of literature about distributed computing exists, it can sometimes be difficult to understand how things can actually be done in practice. Through this, I can help other people to use distributed computing infrastructures for the benefit of their research.

Today, your position at UCL is funded by the CRESTA project though. What's this?

CRESTA has been running for a while, but I only joined the project last year. CRESTA stands for 'Collaborative Research into Exascale Systemware, Tools, and Applications' and participants are working on the optimization of applications for the exascale supercomputing era. This entails work with various partners on issues such as scalability, profiling, monitoring, and visualization. It's all about interaction between application communities and those working hard on the tool-development side.

Finally, why is it that you're interested in attending ISC'14? Why is the event important for early-career researchers such as yourself?

What I find interesting about ISC'14 is that it brings together a very diverse set of HPC communities, not only from Europe, but also from the rest of the world. For me, it's a great way to get informed about what these communities are doing and what people feel are priority topics within HPC. It's a good event for getting ideas for new scientific work and it's also an important place for networking, of course. This is really essential if you're an early-career researcher. You need to find the right people with whom to collaborate, so that you can put your skills to good use and push back the frontiers of knowledge within your field.

And, what are you most looking forward to at the event?

I'm looking forward to finding out about the

other exascale challenges discussed in my session and I hope to go to the exhibition floor to find out more about the emerging supercomputers in Europe. I also plan to speak to lots of people at the event and hope to learn new things about scalability and computational techniques for emerging architectures.

ISC'14will be held in Leipzig, Germany, from 22-26 June, 2014.

Join the conversation

Contribute

Do you have story ideas or something to contribute? Let us know!

FUNDING PARTNERS

CATEGORIES

CONNECT WITH



The National Science Foundation supports

Blood flow in the brain, multi-scale modeling, a...

https://sciencenode.org/feature/blood-flow-brai...

2



the US desk under award 1242759, for sustaining and strengthening International Science Grid This Week (which recently became the Science Node).



CERN, the European Organization for Nuclear Research, supports the Science Node. The organization has played a key role in the publication since 2006, and currently hosts the European editor.

Advanced
computing
Research
networks
Big data
Tech trends
Community
building

\succ
CONTACT
Science Node
Email:
Email: editors@sciencenode
Email: e ditors@sciencenode Website:
Email: editors@sciencenode Website: sciencenode.org

5

Copyright © 2015 Science Node ™ | Privacy Notice | Sitemap

Disclaimer: While Science Node [™] does its best to provide complete and up-to-date information, it does not warrant that the information is error-free and disclaims all liability with respect to results from the use of the information.